Stochastic Phonological Knowledge: The Case of Hungarian Vowel Harmony

Bruce Hayes   Zsuzsa Cziráky Londe

UCLA

August 2005

Abstract

In Hungarian, stems ending in a back vowel plus one or more neutral vowels show unusual behavior: for such stems, the otherwise-general process of vowel harmony is lexically idiosyncratic. Particular stems can take front suffixes, take back suffixes, or vacillate. Yet at a *statistical* level, the patterning among these stems is lawful: in the aggregate, they obey principles that relate the propensity to take back or front harmony to the height of the rightmost vowel and to the number of neutral vowels.

We argue that this patterned statistical variation in the Hungarian lexicon is internalized by native speakers. Our evidence is that they replicate the pattern when they are asked to apply harmony to novel stems in a "wug" test (Berko 1958). Our test results match quantitative data about the Hungarian lexicon, gathered with an automated Web search. We model the speakers' knowledge and intuitions with a grammar constructed under the dual listing/generation model of Zuraw (2000), then show how the constraint rankings of this grammar can be learned by algorithm.*

*We would like to thank Stephen Anderson, Arto Anttila, Andrew Garrett, Matthew Gordon, Gunnar Hansson, Sharon Inkelas, Patricia Keating, Paul Kiparsky, Janet Pierrehumbert, Catherine Ringen, Colin Wilson, Péter Siptár, Donca Steriade, Robert Vago, and Kie Zuraw for helpful advice. As is usual, they are not to be held responsible for defects. We would also like to thank our many Hungarian language consultants for sharing their native-speaker intuitions.

1. Introduction: Irregularity in Phonology

Linguists sometimes have the luxury of working with systematic, exceptionless data. More often, we encounter data which cannot be reduced to a single general pattern. In phonology, the variation is usually lexical: a subset of stems fails to adhere to the most frequent data pattern. This article addresses the question of what the language-learning child does when she confronts such cases. We focus on the appearance of front- vs. back-voweled suffix allomorphs in Hungarian vowel harmony, for example as in *fal-na̲k* 'wall-dative' vs. *kert-ne̲k* 'garden-dative'.

Most previous work concerning irregularity in phonology has adopted an approach in which the majority pattern is characterized as regular, with some mechanism chosen to deal with the residual cases. For instance, in Halle and Mohanan's (1985) account of English past tenses, irregular forms are lexically marked to undergo minor rules. Thus, for instance, the irregular past tense form *clung* is lexically listed as /klɪŋ/ with a special diacritic mark, which causes it to undergo Backing Umlaut, a rule which has the effect of shifting the stem vowel from [ɪ] to [ʌ] in the past tense form.

Another possibility, proposed by Pinker and Prince (1988), is to use grammar to derive only the regular forms, and simply list all the irregulars in the lexicon; thus for Pinker and Prince, *clung* is underlyingly just /klʌŋ/. At first blush such an approach seems inadequate, as without amplification it cannot account for the fact that irregulars usually occur in patterns, such as *cling - clung*, *fling - flung*, *sling - slung*. There is evidence (Bybee and Moder 1983, Prasada and Pinker 1993, Albright and Hayes 2003) that such patterns are partly productive, so a pure-listing account fails to capture the native speaker's knowledge. Therefore, Pinker and Prince amplify their proposal with the idea that the memorized lexical entries for irregulars are embedded in a kind of associative network. This network would be able to generate novel irregulars by some sort of analogy, thus accounting for whatever productivity their patterns may have.

This approach is unsatisfactory to the extent that irregulars can be shown to be derived on the basis of principles of phonological theory, which presumably would be included only in the theory of grammar, and not in the analogical network. Albright and Hayes (2003) have argued that English irregular past tenses are indeed derived by a grammar and not by analogy. The Hungarian data reviewed below arguably form an especially clear case: they are the result of completely ordinary mechanisms of vowel harmony, so it is difficult to justify relegating the minority patterns to a totally different mechanism. Indeed, the minority patterns sometimes compete with the majority on a near-equal basis (see §4.2), which makes it especially arbitrary to claim that the majority results from grammar and the minority from analogy.

We argue here instead for a unified approach to irregularity, as developed by Zuraw (2000). In Zuraw's theory, all of the competing patterns are expressed in a single grammar, along with a characterization of their relative strength. The grammar is implemented in stochastic Optimality Theory (Boersma 1997, Boersma and Hayes 2001). Individual inflected forms, which usually show invariant behavior, are lexically listed where necessary (just as in Pinker and Prince's theory), and their invariance is guaranteed by high-ranking Faithfulness constraints. When a speaker must provide a novel inflected form (for instance, because she has never heard the stem in the relevant inflectional category), the stochastically-ranked constraints of the grammar provide a *range* of options, each with a probability of being output. This probability is determined in the course of language learning, and approximately reflects the frequencies of the competing patterns as they appear in the lexicon. The basic prediction of the model is that the lexical frequencies, insofar as they reflect relevant phonological properties of stems, should give rise to a grammar that generates outputs at frequencies approximating the lexical frequencies.

This is a testable prediction, since we can get speakers to use their grammars to generate novel forms by asking them to inflect stems they have never heard before, in the classical "wug" test paradigm (Berko 1958). Zuraw's own work demonstrates a fairly good match between the frequency patterns of Nasal Mutation in the Tagalog lexicon with the varying intuitions of her wug test subjects, and she develops a stochastic grammar that links the two. Similar work in other frameworks has likewise found a match between statistical patterns in the lexicon and gradient speaker intuition: Eddington (1996, 2004), Albright (2002), Ernestus and Baayen (2003), Albright and Hayes (2003), and Pierrehumbert (in press).

Here, we present a study similar to Zuraw's, focusing on Hungarian vowel harmony. As in earlier studies, we find that native speakers are good frequency matchers: in a wug test, the forms they volunteer closely match the statistical pattern of the Hungarian lexicon.

With this result in hand, we proceed to analysis, with three goals in mind. First, we develop a grammar within Zuraw's framework that accurately describes the wug testing behavior of our consultants. Second, we extend the analysis to an area not addressed by Zuraw, namely the question of *impossible* harmony patterns; i.e. the characterization of patterns that do not exist and (as we will claim) could not exist. Finally, we turn to the question of learnability, suggesting a way in which algorithms from current work can be used to learn the variable Hungarian harmony pattern.

2. Hungarian Vowel Harmony

Hungarian vowel harmony has been the focus of great deal of research which we have freely mined for generalizations and analytic insights. A non-exhaustive list of earlier work includes Esztergár (1971), Ringen (1975), Vago (1974, 1976, 1980), Kontra and Ringen (1986), Hare (1990), Kornai (1991), Ringen and Vago (1995, 1998), Dienes (1997), and Siptár and Törkenczy (2000).

The vowels of Standard Hungarian are given below in (1), first in IPA then in orthography. Since the orthography conveniently depicts natural classes we use it here.[1]

[1] The orthography for consonants matches IPA except in the following cases: *c* [t͡s], *cs* [t͡ʃ], *gy* [ɟ], *s* [ʃ], *sz* [s], *ly* [j], *ny* [ɲ], *ty* [c], *zs* [ʒ].

(1)

	IPA				**Orthography**		
	Front unrounded	Front rounded	Back		Front unrounded	Front rounded	Back
high	i, iː	y, yː	u, uː	high	i, í	ü, ű	u, ú
mid	eː	ø, øː	o, oː	mid	é	ö, ő	o, ó
low	ɛ		ɔ, aː	low	e		a, á

The vowel transcribed as [ɛ] is classified as low, since it is paired with the low vowel *a* in the harmony system.

In what follows, it will be useful to express the vowel sequences of words in formulas, and for this purpose we adopt the following abbreviations: **N** (mnemonic for "neutral") will designate the front unrounded vowels, **F** the front rounded vowels, and **B** the back vowels. For example, in this notation, the word *albérlő* 'lodger' is BNF.

Hungarian is a richly inflected language with dozens of suffixes. We will deal here only with the large class of suffixes that show a two-way alternation in backness; thus we will be ignoring the harmonically invariant suffixes (Vago 1980, 15-18; Siptár and Törkenczy 2000, 65-6), as well as suffixes that show a three-way alternation based on backness and rounding (Vago 1980, 18-19; Siptár and Törkenczy 2000, 72-74). Since the two-way suffixes generally behave alike, it suffices for present purposes to discuss just one of them, namely the dative, which appears as *-nak* or *-nek* according to the principles of vowel harmony.

Vowel harmony depends on the vowels that appear near the end of the stem. For instance, if the last vowel of a stem is back, then no matter what vowels come earlier, the suffix must also be back, as shown in the examples of (2):

(2)	BB	ablak-nak	'window-dat.'
	NB	bíró-nak	'judge-dat'
	FB	glükóz-nak	'glucose-dat.'

Likewise, if the last vowel of a stem is a front rounded vowel, then the suffix vowel must be front:

(3)	F	üst-nek	'cauldron-dat.'
	BF	sofőr-nek	'chauffeur-dat.'

Again, it does not matter what vowels occur earlier in the stem.

Most stems whose vowels are all front unrounded (N) take front suffixes:

(4)	N	kert-nek	'garden-dat.'
	N	cím-nek	'address-dat.'
	NN	repesz-nek	'crack-dat.'

However, there are a few dozen exceptional all-N stems that take back suffixes, even though they contain no back vowels:

(5)	N	híd-nak	'bridge-dat.'
	N	síp-nak	'whistle-dat.'
	NN	derék-nak	'waist-dat.'

Following earlier usage, we will refer to these as *híd* stems, after the word for 'bridge' given in (5) above. All but two are monosyllabic, and of these, most contain the vowel /í/.

The remaining cases are those in which a harmonic vowel (F or B) precedes a string of one or more neutral vowels at the end of the stem. Of these, the stems of the form …FN, …FNN, …FNNN all take front harmony:

(6)	FN	fűszer-nek	'spice-dat'
	FNN	őrizet-nek	'custody-dat.'

The most complex examples, which are the focus of this article, are stems of the type …BN, …BNN, …BNNN, etc. Here, we find extensive lexical idiosyncrasy (Vago 1980, 14, 22; Siptár and Törkenczy (2000, 70-72): individual stems can require back suffixes, or require front suffixes, or allow both in free variation. Thus, for instance, Siptár and Törkenczy cite *haver* 'pal' as a stem that takes only back suffixes (*haver-nak*); *hotel* 'hotel' as a stem that can take either front or back (*hotel-nak ~ hotel-nek*); and *kódex* 'codex' as a stem that only takes front suffixes (*kódex-nek*). A triplet from our own data, in this case with /…Bé/, is the following:

(7)	BN	pallér-nak (only)	'foreman-dative'
	BN	arzén-nak, arzén-nek	'arsenic-dative'
	BBN	mutagén-nek (only)	'mutagen-dative'

A particular speaker of Hungarian can be assumed to memorize for each member of this class of stems which class of suffix allomorphs (front or back) it takes.

3. The Statistical Patterning

While it is not predictable in general whether a BN or BNN stem will take front or back harmony, there are clear tendencies present. If one knows what vowels such a stem contains, it is possible to guess, with far better than chance frequency, what kind of harmony it will take. The crucial generalizations have been studied by Szépe (1958), Vago (1974), Anderson (1980), Kontra and Ringen (1986), Farkas and Beddor (1987), Siptár and Törkenczy (2000), Benus (2005), and other scholars.

The first generalization is what we will call the **height effect**, based on the height of the rightmost vowel in […BN]: the phonologically low vowel *e* (IPA [ɛ]) occurs with front suffixes more often (that is, in proportionately more stems) than the mid vowel é ([eː]), which occurs with front suffixes more often than the high vowels *i* and *í* ([i, iː]). Second, there is a **count effect**: BNN stems take front suffixes more often than BN stems do.

Since the sources cited above do not agree on the precise nature of the height and count effects, we have sought to collect as many data as possible on these quantitative patterns. We have followed two methods: elicitation from native speakers of large number of forms, and a machine-based search of the World Wide Web. Since the latter has turned up more data, we will cover it first.

3.1 A search-engine study of the Hungarian lexicon

The basic method of collecting quantitative patterns for phonology by using a Web search engine was pioneered by Zuraw (2000). The idea is that where forms occur in free variation, we can measure their relative frequencies by counting the hits returned for each.

To see how this works, consider the forms from (7) above. A query for these forms using the Google search engine (12 May 2004) yielded the hit counts shown in (8).

(8)

Word	Hits	Percent
mutagén-nak	0	0%
mutagén-nek	32	100%
arzén-nak	32	67%
arzén-nek	12	33%
pallér-nak	15	100%
pallér-nek	0	0%

These hit counts agree with findings obtained by casual elicitation from native speakers; namely that *mutagén* takes front suffixes, *pallér* takes back suffixes, and *arzén* can take either. In other words, native speaker intuition matches native speaker behavior, i.e. the behavior of Hungarian speakers who happen to be using the dative form of these stems when composing a Web page.

In our study, we searched not just a few representative stems, but a long list (10,974 stems), taken from a Hungarian electronic lexicon. We did this with a computer program that queried the search engine automatically. The forms that were fed to the program were constructed by adding -*nak* and -*nek* to each stem and applying the rule of Low Vowel Lengthening (Vago 1980, 3-4; Siptár and Törkenczy 2000, 56-58) where appropriate.[2] We kept data for words in which the search yielded ten or more total hits (-*nak* and -*nek* summed).

Since obtaining phonological data from the Web is a fairly new technique, we mention here a few precautions. First, it should be remembered that a search engine does not count actual tokens of the target form, but only the number of Web pages that contain it. We believe that where the focus of interest is *relative* frequency (here, of front vs. back endings), this factor will

[2] Our software, called "Query Google," was programmed by Timothy Ma of UCLA. It is implemented as a publicly accessible Web applet; address http://www.linguistics.ucla.edu/people/hayes/QueryGoogle/. The electronic dictionary from which we obtained our words is at http://sourceforge.net/project/showfiles.php?group_id=41308&package_id=33412.

impose only minimal distortion, particularly where the stems under investigation are not especially common.

Second, at least some of the data retrieved in a Web search will be nonsensical in some way. For instance, speakers who lack Hungarian keyboards occasionally leave off umlauts or acute accents, which can distort the results for generalizations based on backness or length. Compound words always take the harmony of the second member, which can create errors if they are mistakenly counted as monomorphemic. Borrowed words are occasionally spelled as in the source language but take harmony according to their Hungarian pronunciation; thus, for example, *Birmingham* is pronounced [ˈbörmingem] and thus takes front harmony, despite its orthographic *a*.

To ward off trouble from these sources, we did some hand checking. We went through every instance of BN and BNN stems in the corpus, eliminating the illegitimate examples. We also checked all the baffling instances like B stems taking front harmony or F or FN stems taking back, and found that (with just a tiny residue of completely mysterious forms), the exceptions could be accounted for on the basis of the above categories.

As always in such studies, we must consider whether to count **tokens** (e.g. "365,822 occurrences of BN stems in the corpus take -*nak*") or **types** (e.g. "603 BN stems in the corpus take -*nak*"). The literature suggests that when the extension of morphological patterns in the lexicon is at stake, it is type frequency that is primarily relevant; for discussion see Bybee (1995, 2001), Pierrehumbert (2001), Albright (2002), and Albright and Hayes (2003). Our findings are in harmony with these earlier claims, as we found that types provide a better fit to native speaker judgment (see §4.2). We will therefore report only type frequencies here.

In counting type frequencies we assigned vacillators to the front and back categories according to the percentage breakdown of each type; thus for a vacillator that took -*nak* 20% of the time and -*nek* 80% of the time, we would add .2 to the total of -*nak* stems in its category and .8 to the total of -*nek* stems.

We report our data with a "backness index": for any particular phonological category (such as BN), the backness index is the proportion of stems in that category that took -*nak*, counting vacillators as just noted. The backness index for a category takes the value 1 when every stem always takes -*nak* and 0 where every stem always takes -*nek*.

Lastly, to help relate our findings to previous work, we also include a sorting into "back", "front" and "vacillator" forms, where "back" is arbitrarily defined as taking back suffixes at least 97% of the time, "front" as taking front suffixes at least 97% of the time, and vacillatory any other form.

(9) Findings of the Google Survey

Stem Type	Back	Vacillator	Front	Total stems	Backness Index
B	6251	39	0	6290	.999
BN	603	78	83	764	.831
Bi	458	17	0	475	.989
Bí	52	0	1	53	.980
Bé	93	18	9	120	.845
Be	0	43	73	116	.104
BNN	6	21	44	71	.206
BNi	1	12	17	30	.223
BNí	1	7	0	8	.358
BNé	4	2	6	12	.421
BNe	0	0	21	21	0
N	14	23	259	296	.078
NN	0	4	929	933	.002
F	0	0	698	698	0

The Google survey strongly confirmed the generalizations stated in §3 above. The chart in (10) provides a coarse classification of our data, ignoring vowel height for the moment:

(10) Google Data: Basic Stem Types

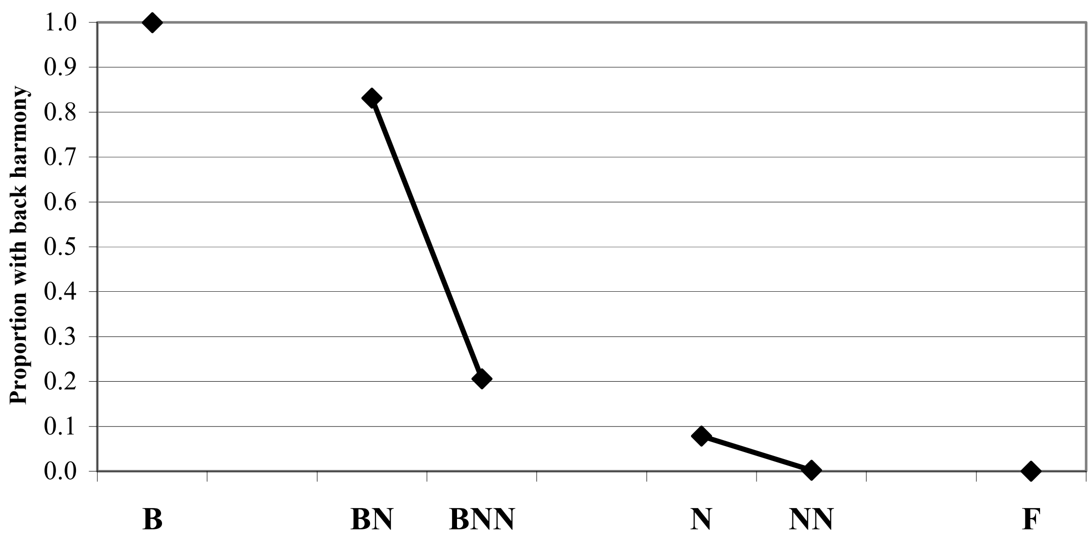

As can be seen, stems ending in F always take front harmony (that is, the backness index for the 698 stems examined was zero). Stems ending in B virtually always take back harmony (.999; we assume that the exceptions were typographical errors). Stems with all neutral vowels (N, NN) are occasionally *híd* stems when monosyllabic (.078) and only rarely when disyllabic

(.002). Crucially, comparing the BN vs. BNN stems, we find a strong confirmation for the count effect in the much lower backness index for BNN (.232 overall) vs. BN (.840).[3]

Turning to the height effect, we first sort the relevant forms (all BN and BNN) by the height of their last vowel: Bi, Bí, BNi, and BNí all have high rightmost vowels, and form the High category; Bé and BNé all have mid rightmost vowels, and form the Mid category; and Be and BNe all have low rightmost vowels, and form the Low category.

(11) Google Data: Height Effect

Stem Type	**Back**	**Vacillator**	**Front**	**Total stems**	**Backness Index**
High	512	36	18	566	.938
Mid	97	20	15	132	.806
Low	0	43	94	137	.089

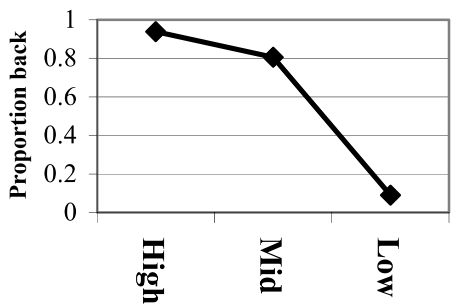

Both the High-Mid and the Mid-Low differences are highly significant; see fn. 3.

Lastly, we must consider how the two effects interact: do both BN and BNN forms have a height effect? The data here are equivocal:[4]

(12) Google Data: Height Effect in BN vs. BNN

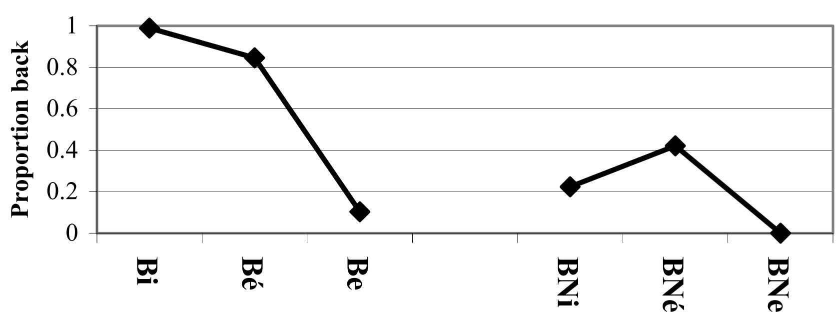

[3] We submitted the data to chi-square tests, dividing the vacillators between front and back in the same way described above. For BN vs. BNN, $\chi^2 = 146.856$, $p < .001$; for NN vs. N, $\chi^2 = 65.248$, $p < .001$. For the height effect described in the next paragraph: High vs. Mid, $\chi^2 = 23.489$, $p < .001$; Mid vs. Low, $\chi^2 = 140.205$, $p < .001$.

[4] Looking ahead to comparison with our experimental data, we omit forms with /i/, which are very few.

The height effect is clearly evident in the BN forms, which are numerically preponderant. However, the numbers for BNé and BNi are surprisingly reversed with respect to Bé vs. Bi. This fact will be relevant below when we consider the preferences of Hungarian speakers for novel forms.

Lastly, it appears that when both the height and count effects are maximally in effect, the lexicon is variation-free: as Siptár and Törkenczy (2000, 71) point out, all BNe stems take front suffixes.

3.2 Verification with native speaker consultants

As a check on the search engine method, we selected from our lexical list all instances of BN, BNN, and N, plus representative cases of NN, for a total of 1130 stems, and asked two adult native speakers of Hungarian from Budapest to indicate for each whether they preferred to use -*nak*, preferred -*nek*, or could use either. The speakers did not rate exactly the same stems as the Google survey, because a few of the Google words were unfamiliar to them, and the speakers rated a number of words that failed to reach our threshold of ten hits on the Google survey. The overlapping forms totaled 769 for Speaker 1 and 767 for Speaker 2.

The degree of agreement is given in Table (13). In the table, "Vacillators" for Google data are defined as in the previous section, and for the native speakers are simply the stems for which they volunteered both options. "Off by one" means one source gave -*nak* or -*nek* and the other gave "vacillator".

(13)

		Google: -*nak*	**Google: vacillator**	**Google: -*nek***			
Speaker 1	***-nak***	272	32	3	*Agree*	676	87.9%
	vacillator	1	22	12	*Off by one*	90	11.7%
	-nek	0	45	382	*Disagree*	3	0.4%
Speaker 2	***-nak***	269	23	0	*Agree*	676	88.1%
	either	4	46	34	*Off by one*	91	11.9%
	-nek	0	30	361	*Disagree*	0	0%

The agreement seems fairly clear. Further, if we assign numerical values to the speakers' responses (-*nak* = 1, vacillator = .5, -*nek* = 0) and compute the correlation coefficient with the Google backness indices, we find r = .951 for Speaker 1 and .937 for Speaker 2, slightly better than the speakers' agreement with each other (r = .915). We conclude that the search engine method gives results quite similar to those provided by individual native speaker consultants. We are inclined, in fact, to trust our Google data better: they are based on far more data, and they represent people actually using their language rather than making metalinguistic judgments.

4. The Productivity of the Pattern: A Wug Test

Are the height and count effects mere statistical patterns of the Hungarian lexicon, or are they actually internalized by Hungarian speakers and extended productively? The usual test for answering this question is the "wug" test, pioneered by Berko (1958), in which productivity is assessed by asking speakers to inflect novel stems. In the wug test we conducted, we gave speakers new, made-up stems in the nominative case (that is, with no suffix), and set up the experiment to elicit these stems with the dative suffix, either *-nak* or *-nek* as the subject chose. Our experiment extends and complements work by Kontra and Ringen (1986) and Gósy (1989), who tested loan words.

4.1 Procedure

We chose our wug stems on the basis of several criteria. First, we included both BN and BNN stems, with at least one stem in each category ending in each of the vowels /i, é, e/. In order to sample the rest of the stem inventory, we included stems ending in F and B and as well as monosyllabic and disyllabic neutral-voweled stems. We made two such sets of 15 wug stems; any particular consultant saw just one of the two sets, chosen at random. The two sets were as shown in (14).

(14)

Type	**Set 1**	**Set 2**
Bi	monyil, csádik	kánit, pozin
Bé	hádél, kolén	vuszék, vánél
Be	órel, bontel, kázen	ranyel, unyeg, csúltek
BNi	poribit	lolivit
BNé	lányitég	ányivél
BNe	fányedeg, luteker	álendel, móleter
N	híny	nyís
NN	zefét	petlér
F	gyülüt	hösög
B	szandat	bortog

In constructing these wug stems, we attempted to make them sound as phonologically ordinary in Hungarian as possible. This was done by extracting from our electronic dictionary the most common initial, medial, and final consonants and consonant clusters, and using these to provide the consonants. We also attempted to avoid wug stems that had a strong direct resemblance to any particular existing stem, or would be likely to be interpreted as compounds.[5] This was done by generating large numbers of candidates for each type and checking them according to the native intuition of the second author and several other native speakers.

[5] This turned out to be harder than we thought: Péter Siptár has pointed out to us the resemblances *vuszék ~ szék* 'chair', *ányivél ~ vél* 'think, opine', *ranyel ~ nyel* 'swallow', and *monyil ~ nyíl* 'arrow'. Future work in this area should probably check for such resemblances by machine as well as by eye.

We administered the wug test as a written questionnaire.[6] The wug words were given in paragraph frames, meant to give the participants practice in using them before constructing their dative forms. The first sentence of the paragraph provided the nominative form of the wug stem, the second required the consultant to repeat the nominative by filling in a blank, and the third provided a grammatical context requiring the consultant to use the dative case. Frames and instructions were composed with the goal of encouraging the subjects to treat the stems as long-forgotten but authentic words of Hungarian, rather than as recent loans. Here is an example of one of our frames translated into English; italics indicate expected responses:

(15) Sample wug test frame

> **hádél**
>
> Women in the Middle Ages used hádél to wash clothing. Back then, __(*hádél*)__ grew abundantly in the fields. It is very hard to find nowadays, but it is said that __(*hádélnak* or *hádélnek*)__ had a wonderful fragrance.

We used multiple versions of each test, to make sure that no wug stem was consistently affiliated with a particular frame, and we also changed the order of the wug stems and frame sentences at random in the various versions.

The test was distributed to 171 Hungarians in Budapest (161 consultants) and Tiszafüred (10). The consultants, who were mostly young adults, took the test as a favor to us.

4.2 Results

Since wug stems of the same type (e.g. *órel*, *bontel* are both /Be/) received quite similar scores, we pooled their scores into single categories, obtaining the results shown in (16) below. For comparison, this chart also includes the Google data from the previous section, shown in

[6] As it turned out, use of written presentation provided an important additional control, since there is now some evidence that there are tiny phonetic differences between Hungarian stems that have the same basic vowels but take different harmony; see Benus (2005), Benus and Gafos (in press). Hungarian orthography provided the subjects with an unbiased characterization of the gross vowel phonemes present, without the possible interference of microphonetic distinctions.

gray.[7]

(16) Wug Test Results Compared with Google Survey Results

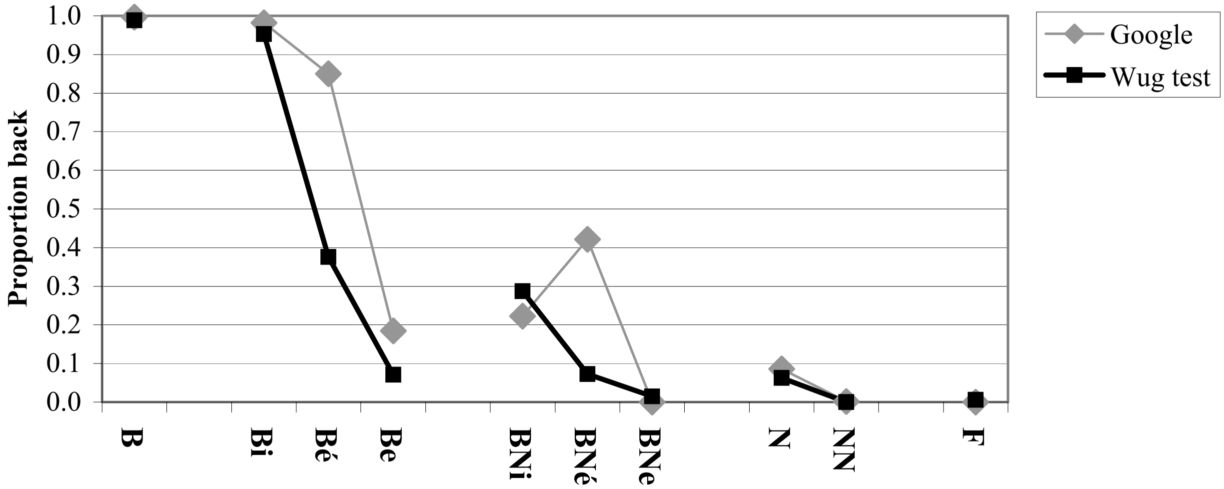

Examining the individual cases, we can see that the experiment yielded sensible results for areas where the suffix choice is obvious: virtually all instances of F stems took front harmony, and virtually all instances of B stems took back. Thus our consultants appear to have understood the task and performed reliably.

In stem types for which the Hungarian lexicon includes both front and back cases, the aggregate behavior of the consultants tended to statistically match the proportions found in the lexicon. For instance, according to our data about 7.8% of the monosyllabic N stems in the Hungarian lexicon are *híd* stems, taking back harmony. In the wug experiment, 6.3% of our consultants interpreted *híny* and *nyís* as if they were *híd* stems, attaching -*nak*. NN stems are very seldom of the *híd* type in the lexicon (0.2% in the Google survey), and none of our consultants produced back harmony for either of our NN stems *zefét* and *petlér*.

[7] These are the raw values for the subject data in (16). The invalid responses were mostly blanks and forms given with no suffix.

Stem Type	Back	Front	Invalid response	Backness Index
B	165	2	4	0.988
Bi	322	16	4	0.953
Bé	127	211	4	0.376
Be	35	456	22	0.071
BNi	48	119	4	0.287
BNé	12	153	6	0.073
BNe	5	330	7	0.015
N	10	148	13	0.063
NN	0	171	0	0.000
F	1	164	6	0.006

The height and count effects found in the Google survey for BN and BNN stems also emerged in the subject responses for the wug experiment. As can be seen in (16), the lower the final stem vowel, the more front responses we obtained; and more front responses were obtained for BNN than BN; chi-square testing showed these differences to be highly significant.[8]

Lastly, we assessed the overall degree of agreement between the Google survey and the wug test results. To do this, we took each of the 30 stems tested (see (14)), and paired it with the backness value obtained in the Google survey for its general category (such as Bi, BNi, etc., as in (16)). The correlation found was r = .896, indicating fairly close agreement. We note in passing that if the proportions in the Google data are calculated from token rather than type frequencies, this correlation emerges as somewhat lower (r = .820); cf. discussion above in §3.1.

4.3 Smoothing

Some of the discrepancies between the Google survey and the wug test data seem of potential importance. In the wug data, there is an across-the-board height effect even in the BNN forms: BNi stems took more back suffixes than BNé stems, which in turn took more back suffixes than BNe stems. The data from the Google survey contradicted this pattern, with more back responses for BNé than for BNi.

In our view, it is the Google data that most likely are aberrant. At this level of phonological detail, there are only a few relevant stems in the lexicon. The figure of .421 found for BNé is based on just 4 back stems, 2 vacillators, and 6 front stems, for a total of 12.

What is interesting is that if the aberrant figure of .421 does represent the Hungarian lexicon as a whole, the aberrance is evidently not registered by native speakers, whose wug test values for BNN indicate a straightforward height effect, with higher vowels taking more back suffixes.[9] We conjecture that the speakers have in some sense **smoothed** the data. Rather than reflecting every small idiosyncrasy in the Hungarian lexicon, they formulate more general patterns based on natural phonological dimensions, namely the height effect and the count effect.

The surprising 1.5% of cases where our consultants volunteered *-nak* for BNe stems, contradicting the unanimous lexical pattern, are plausibly also an instance of smoothing. The fact that BNN stems in general can take *-nak*, and B…e stems in general can take *-nak*, may have led our consultants to arrive at the marginal possibility that BNe stems, which intersect these two categories, can take *-nak*.

---

[8] For BN vs. BNN, $\chi^2 = 203.716$, $p < .001$; for High vs. Mid, $\chi^2 = 209.924$, $p < .001$; for Mid vs. Low, $\chi^2 = 139.338$, $p < .001$.

[9] Wug test: BNi vs. BNé, $\chi^2 = 25.839$, $p < .001$; BNé vs. BNe, $\chi^2 = 11.246$, $p < .001$. Google survey: BNi vs. BNé, $\chi^2 = 0.83$, not significant; BNé vs. BNe, $\chi^2 = 7.445$, $p = 0.0064$. For both Google survey comparisons, Yates's correction for small values was employed.

4.4 Caveats

Before continuing with a formal analysis of our data, we discuss a possible confound and a puzzle.

First, it has been suggested to us that our wug test results merely reflected a mixture of differing idiolects. For example, the .376 backness value for Bé stems could have resulted from 37.6% of the speakers having an idiolect that always assigns back endings to Bé stems, and 62.4% having an idiolect that always assigns front endings. We checked this hypothesis by testing cases where the very same consultant rated two different stems with the same pattern. For instance, *hádél* and *kolén*, both Bé, appeared on the same questionnaires, and thus were encountered in the same session (though not consecutively) by the same speakers. In a series of chi-square tests, we found that consultants who gave back responses for *hádél* were no more likely to give back responses for *kolén* than consultants who gave front responses for *hádél*. We obtained similar results for all other pairs where enough data were available for testing. We conclude that idiolect variation played at most a minor role in our results. The level of variation is not between individuals, but within the individual: when confronted with a novel wug stem, each speaker behaved stochastically, in a way that matched the frequencies of the lexicon.

There is one respect not yet discussed in which the wug test results diverged from the lexicon: overall, in comparison to the Google survey, the wug test subjects preferred front suffixes; averaging across the ten categories given in chart (16), the wug test values are .091 more front. Concerning this puzzle we offer the following conjecture. For reasons we do not understand, there is a weak connection in the Hungarian lexicon between stem frequency and frontness: the rarer the stem, the more likely it is to take front endings. Thus, we find that within just the BN stems, the correlation of the frequency of the bare stem (which we also measured in the Google survey) with the proportion of back responses is r = .141.[10] Wug stems are, by definition, the rarest of stems (frequency zero), and this may have contributed to their slightly front-preferring behavior. In principle, this factor could be entered into the model described below, but we will not attempt to do this here.

5. A Theoretical Model of Variation in Hungarian Vowel Harmony

We turn to the task of developing an explicit analysis of our findings, drawing on various notions from current phonological theory. To achieve descriptive adequacy our model must accomplish three tasks.

First, it should accommodate **stem-specific behavior**, permitting speakers to list (in some form yet to be addressed) what kind of suffixes, back or front, are taken by a particular BN or BNN stem (cf. (7) above). Most Hungarian BN and BNN stems are *not* vacillators, and learning stem-specific behavior is part of what is involved in learning the Hungarian lexicon.

[10] This connection might explain Gósy's (1989) finding that younger children tend to give backer responses in wug testing for the relevant word classes. Younger children would be less likely to be familiar with rarer words.

Second, an adequate model should characterize the native speaker's **expectations** about what suffixes a novel stem will take—in particular, it should be able to account for our wug test data.

Third, a model should characterize the **limits of stem-specific behavior**. While it is true that BN and BNN stems can have their own specific behavior, B or F stems cannot; their harmony pattern is completely predictable. For example, our wug test included B stems like *bortog* and *szandat*, and these virtually always took back harmony; likewise for F stems like *gyülüt* and *hösög* taking front harmony. We claim that forms like **bortog-nek* or **gyülüt-nak* are simply unacceptable in Hungarian, and that a phonological analysis should capture this fact.

5.1 Theory

We assume Optimality Theory (Prince and Smolensky 1993/2004), in which the outcomes of phonological derivations depend on the ranking of conflicting constraints. Constraint conflict arises here when a stem contains both front and back vowels. A variety of constraints require that the suffix vowels match the stem vowels in backness, and when a stem contains both front and back vowels, these constraints will conflict.

We use a **stochastic** variant of Optimality Theory (Boersma 1997, Hayes and MacEachern 1998, Boersma and Hayes 2001), in which the ranking of constraints is probabilistic: every constraint pair (A, B) is associated with a probability (0-1) specifying how likely it is that A will dominate B on any given speaking occasion. The reason for using stochastic OT is that it permits precise predictions about the relative proportions of forms produced in free variation.

Lastly, we assume the **dual listing/generation model** of Zuraw (2000). In this model, grammars may contain sets of Markedness constraints that are stochastically ranked with respect to each other, but subordinated to Faithfulness constraints. This means that existing forms, which are covered by a particular lexical entry and protected by Faithfulness, surface without variation; whereas newly inflected forms, where Faithfulness constraints are inapplicable, are derived stochastically according to the pattern of the subordinated Markedness constraints.

The following sections cover underlying forms, constraints and rankings, and an assessment of the model's performance in describing our data.

5.2 Underlying forms

The underlying phonological form of Hungarian suffixes is somewhat difficult to establish, since they normally appear harmonized to a preceding vowel. The occurrence of case suffixes as independent stems, in constructions like *nek-em* 'me-dat.', could in principle justify an underlying backness value for these suffixes (Vago 1973). However, we will see later on that there is reason to let both allomorphs of a harmonizing suffix serve as underlying forms, and so we will assume here that the use of the case suffixes as stems represents an arbitrary lexical choice and does not determine a unique underlying form. For discussion, see Reiss (2003). Below, in cases where it is not crucial to assert an underlying backness value, we will simply use capital letters to designate the general category of the suffix vowel; thus /A/ is a vowel that alternates between [a] and [e], so the underlying form of **-***nak* ~ **-***nek* will be shown as /-nAk/.

5.3 Markedness constraints governing harmony

We assume that BN and BNN stems vary in their harmony because they have two triggers: one which is strong but nonlocal (B), and another which is weak but local (the rightmost N, which is closest to the suffix). Suffix variation results from stochastic ranking of the conflicting constraints that require suffixes to agree in backness with these triggers. The constraints are assumed to be members of the AGREE family (Lombardi 1999, Kiparsky and Pajasalu 2003), relativized to distance.

In our usage, a vowel harmony constraint is **local** if it assigns violations to vowel sequences that are separated only by a (possibly null) consonant string. A constraint is **distal** if it assigns violations without regard to intervening material. In formalizing the constraints, we assume that phonology makes available a vowel projection (Vergnaud and Halle 1979) or tier (Archangeli and Pulleyblank 1987, Clements and Hume 1995) which expresses just the vowels of the string, so that consonants can be ignored in the structural description of constraints. Thus the constraints we will call LOCAL B and DISTAL B are stated below:

(17)a. LOCAL B

*[+back][–back]

Assess a violation when the closest vowel following a [+back] vowel is [–back].

b. DISTAL B

*[+back] X [–back]

Assess a violation when a [+back] vowel is followed somewhere in the word by a [–back] vowel. [11]

For example, in the candidate form (18) below, DISTAL B incurs four violations, as shown by the arrows.

(18) m **u** t **a** g **é** n n **e** k *mutagén-nek*

In contrast, LOCAL B assesses violations only for vowel pairs separated by (at most) a consonant string, e.g. just one violation for *mutagén-nek*:

[11] For earlier analyses that also assume nonlocal harmony, see Kiparsky and Pajusalu (2003), using constraints; and Esztergár (1971, 29) and Vago (1976, 252), using rules.

(19) m **u** t **a** g **é** n n **e** k

We assume additional agreement constraints defined on particular vowels or natural classes of vowels, formalized analogously to (17a,b). In the present analysis the following constraints will be employed:

(20) a. Local F, where F = [–back, +round]
b. Distal F
c. Local i
d. Local é
e. Local e

In principle, Local í should also be included, but we will ignore it here since we have no wug test data for this vowel (in the Google data, it matches /i/ fairly closely.)

We also need a constraint to enforce the more frequent appearance of front suffixes in BNN stems. Walker (2001) has noted that it is possible for harmony processes to have "double triggers"; i.e. harmony occurs only when two in a row of the relevant triggering class are present. Walker analyzes the phenomenon in depth; for present purposes we will just stipulate a constraint Local NN, violated when NN is followed by B.

(21) Local NN

*[–back][–back][+back]

In the suffixed forms of BN stems, Distal B conflicts with one or more of the three constraints Local i, Local é, Local e; in BNN stems, it additionally conflicts with Local NN.[12]

5.4 The height effect in BN stems

With constraints (20c-e) in place, we can propose a tentative explanation for the height effect. In a BN stem, B and N are conflicting harmony triggers. We propose that the propensity of such stems to take front harmony depends on the "strength" of N as a trigger for front harmony.

Kaun (1995, 2004), in a study of the typology of rounding harmony, proposes an explanation for what makes a harmony trigger strong. For Kaun, the differences in the strength of vowels as harmony triggers depend on the phonetic salience with which these vowels manifest the harmonic feature: harmony is triggered preferentially by perceptually inferior vowels, that is, those vowels that lack the extreme phonetic realization of their category. In the case of rounding harmony, these are the low rounded vowels, which (relative to their high counterparts) are

[12] An alternative to Local NN is to fragment Distal B, splitting it into constraints requiring agreement with B two syllables away (governing BN forms) vs. three (governing BNN). We have explored this kind of analysis and find it works about as well as the one presented in the text. For reasons of space we will not present both analyses here.

phonetically less rounded and acoustically less distinct from unrounded vowels. In a typological survey, Kaun found that low rounded vowels often trigger harmony in contexts where high rounded vowels do not. Her functional explanation for this tendency is that the low rounded vowels, which most need help in identification, are more likely to obtain this help by spreading their rounding feature across the word.

Pursuing the same approach for backness harmony, we note that of the front vowels of Hungarian discussed here, it is the lower front vowels that have the lowest second formant frequencies, and thus are perceptually inferior relative to the higher front vowels. Following Kaun's approach, we expect the strength of the front triggers to be determined by their height, with /e/ (phonetically [ɛ]) the best trigger, /é/ ([eː]) the second best, and /i/ the worst. In grammatical terms, this is manifested in an a priori ranking preference:

(22) LOCAL e >> LOCAL é >> LOCAL i

It will be seen below that this ranking, in a looser stochastic form, is indeed what is needed for the analysis of the Hungarian data.[13]

For suggestions that the height effect may not be unique to Hungarian (and thus deserves a general explanation), see Esztergár (1971) and Anderson (1980); for a different phonetic account, see Benus (2005), Benus and Gafos (in press).

5.5 Faithfulness constraints

Following Ringen and Vago (1998), we assume two Faithfulness constraints governing backness; one is limited to root vowels, while the other is simply the general IDENT constraint for this feature:

(23) IDENT-IO(back)$_{root}$

> Assess a violation if a vowel belonging to a morphological root differs in surface representation from its underlying correspondent in its value for the feature [back].

(24) IDENT-IO(back)

> Assess a violation if a vowel differs in surface representation from its underlying correspondent in its value for the feature [back].

As we will see, IDENT-IO(back)$_{root}$ must be ranked higher than IDENT-IO(back), reflecting the greater immutability of root vowels relative to suffix vowels in Hungarian. The pattern of greater faithfulness in roots is often observed cross-linguistically; see for example Casali (1997).

[13] Constraint families based on phonetic scales like (22) have been implemented in various ways. The approach in (22) affiliates one constraint with each member of the scale, and ranks the constraints so as to match the scale. Another approach (Prince 1997, DeLacy 2004) implements the scale by stating the constraints with cutoff points, for example, "Agree in backness if the trigger is mid *or lower*". We have implemented our analysis under both approaches and achieved equally good matches to the data. For brevity we only report the first approach here.

5.6 Constraint rankings: strict

We can now consider how the constraints given above can be ranked to characterize the data. We begin with some straightforward rankings that are nonstochastic (in the theory assumed, they are associated with the probability value 1).

To begin, IDENT-IO(back)$_{root}$ must be ranked strictly over LOCAL B and DISTAL B. This is because Hungarian stems, unlike suffixes, are not in general required to respect harmony: there are many BN and NB stems, and also a fair number of borrowings like *glükóz* 'glucose', with FB, and *sofőr* 'chauffeur', with BF. The need for a strict ranking is shown below in tableau (25): /farmer/ 'blue jeans' survives intact, despite its violations of the two harmony constraints.[14]

(25)

/farmer/	IDENT-IO(back)$_{root}$	LOCAL B	DISTAL B
☞ farmer		*	*
*farmar	*!		

There are also strict rankings among the Markedness constraints. Thus, LOCAL B must strictly dominate DISTAL F, because in all stems in which the last vowel is B and some other vowel is F, the suffix must surface as back—the local B vowel wins out as trigger over the distal F vowel. This can be seen in the following tableau for *glükóz-nak* 'glucose-dat.'.

(26)

/glükóz-nAk/	IDENT-IO(back)$_{root}$	LOCAL B	DISTAL F
☞ glükóz-nak			**
*glükóz-nek		*!	*
*glükőz-nek, *glukóz-nak	*!		

For the same reason, LOCAL F must strictly dominate DISTAL B, for example to obtain *sofőr-nek* 'chauffeur-dat.', not **sofőr-nak*:

(27)

/sofőr-nAk/	IDENT-IO(back)$_{root}$	LOCAL F	DISTAL B
☞ sofőr-nek			**
*sofőr-nak		*!	*
*sofór-nak, *söfőr-nek	*!		

5.7 Constraint rankings: stochastic

In the model of stochastic OT adopted here, the probabilistic rankings of the constraints are expressed by assigning them values along a numerical scale of "ranking strength"; from this scale the relative ranking probabilities can be deduced with a standard mathematical formula,

[14] There is more to the problem than this, in that B combines with N freely in stems, but far less often with F; the BF and FB stems are all borrowings and are felt to be foreign. While we will not try to integrate this fact into the analysis, we believe the necessary apparatus is at hand: Kiparsky and Pajusalu (2003) propose constraints that penalize BF/FB but not BN/NB, and McCarthy and Wolf (2005) offer an update on the concept of the null parse candidate, plausibly a stochastic competitor with BF and FB stems.

given in Boersma (1997, 45). In the discussion that follows, we will present the pairwise probabilities, since these are more readily interpretable.[15]

Three crucial probabilistic rankings in the analysis are those of DISTAL B against its competitors among the weak front-harmony triggers, namely LOCAL e, LOCAL é, and LOCAL i. We propose that these ranking probabilities should be as shown in (28):

(28)

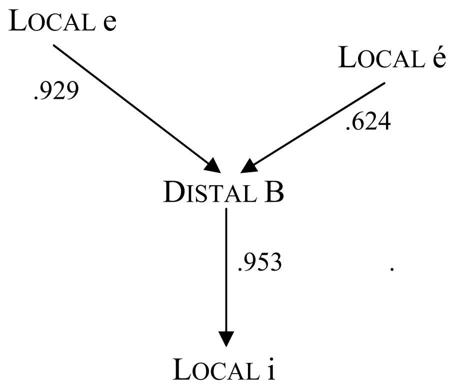

The .624 probability proposed for the ranking of LOCAL é over DISTAL B means that given an input like /hádél-nAk/ (*hádél* is the wug stem appearing in (15)), there is a probability of .624 that the grammar will output *hádél-nek*. This is shown in tableaux (29) and (30).

(29)

.624

/hádél-nAk/	IDENT-IO(back)$_{root}$	LOCAL B	LOCAL e	LOCAL é	DISTAL B	LOCAL i
☞ (.624) hádél-nek		*			**	
hádél-nak		*		*!	*	
hádál-nak	*!					

In (29), the probability of .624 that LOCAL é will dominate DISTAL B implies the same probability that *hádél-nek* will be output by the grammar. The opposite ranking, generating *hádél-nak* with a probability .376, is given in (30):

[15] One set of values that yields the probabilities proposed in this section is: LOCAL e = 105.637, LOCAL NN = 103.285, DISTAL B = 101.483, LOCAL é = 102.377, LOCAL i = 96.746. Strict rankings, such as those from the previous section, can be implemented with any difference in ranking value (e.g., 20) that translates into something very close to 1.

(30)

.376

/hádél-nAk /	Ident-IO(back)$_{root}$	Local B	Local e	Distal B	Local é	Local i
hádél-nek		*		**!		
☞ (.376) hádél-nak		*		*	*	
hádál-nak	*!					

Thus, over a large number of trials, we would expect *hádél-nek* to be the winner in about 62.4% of the trials, and *hádél-nak* to be the winner in about 37.6%. In fact, in our wug test, with stems of the Bé type, this was the percentage obtained from the participants as a whole; our hypothesized ranking values were set up with the express purpose of mimicking this frequency. The remaining values in (28) similarly can be used to derive the correct wug test percentages for Bi and Be stems: the probability .953 that Distal B >> Local i predicts 95.3% back suffixes for Bi stems, and the probability .929 that Local e >> Distal B predicts 92.9% front suffixes for Be stems.

The results so far merely indicate that the constraint set is sufficiently rich to discriminate between the Bi, Bé, and Be categories—we have set three ranking probabilities, and have derived three relative frequencies. More interesting is the task of extending the analysis to the BNN stems, specifically BNi, BNé, and BNe. In the wug test data, these stems exhibited both the height effect and the count effect; the two effects are additive in the sense that the stems that are most likely to take front endings are the BNN stems with low final vowels. We propose that this can be modeled simply by assigning an appropriate probability to the ranking Local NN >> Distal B: this will shift the percentages of back suffixes downward in BNN stems relative to analogous BN stems, because two constraints rather than one are working against Local B.

We have calculated that a probability of .738 for Local NN >> Distal B best fits the data. The crucial part of the grammar is thus as in (31):

(31)

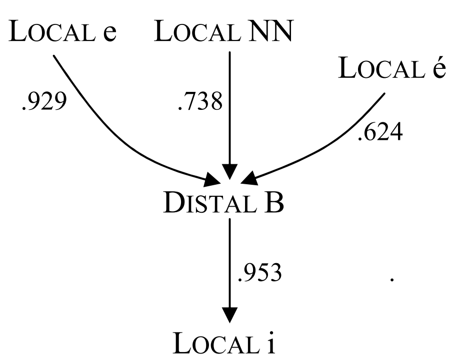

Tableau (32) illustrates how the constraints interact in the case of a representative BNé stem:

(32)

.738
.624

/ányivél-nAk/	Local e	Local NN	Local é	Distal B	Local i
☞ (.836) ányivél-nek				***(!)	
☞ (.164) ányivél-nak		*(!)	*(!)	**	

In cases like this where three constraints stochastically interact, the probabilities can be most conveniently calculated empirically, by running the grammar many times. This procedure when run repeatedly produces very consistent results, and thus is safe to use here.[16]

Testing the grammar in this way for the six crucial cases yields the predictions given in (33).

(33)

	Wug test	Model		Wug test	Model
Bi	.953	.953	BNi	.287	.257
Bé	.376	.376	BNé	.073	.164
Be	.071	.071	BNe	.015	.045

The grammar does not achieve an exact quantitative match for the BNN forms, but it is not far off, and moreover it captures the correct qualitative generalizations: there is a count effect, a height effect, and an additive effect as well—the lowest number of back suffixes occurs for BNe stems, which maximally invoke both effects. This illustrates the ability of the stochastic OT framework to capture additive effects.

The grammar also accomplishes the two instances of "smoothing" we noted in §4.3. Although in the Hungarian lexicon, BNé unexpectedly takes *-nak* more often than BNi, our model in fact gives BNé-*nak* a lower frequency than BNi-*nak*—just as our wug test participants did. Moreover, the wug test subjects unexpectedly volunteered a small number of BNe-*nak* forms, despite their absence in real Hungarian. Our model likewise generates a small number of these forms. In both cases, the cause of the smoothing is the same: the constraints responsible for the height effect are ranked on the basis of all of the data, not just the BNN forms. The statistical patterns found in the (more numerous) BN cases are carried over to some extent to BNN.

5.8 Treatment of existing stems: the role of Faithfulness

Recall that the Hungarian variation is primarily stem-by-stem variation and not token-by-token variation; only the vacillators actually permit the two outcomes generated by the grammar thus far, while most stems impose an invariant suffix choice. This information must be encoded in the lexicon: part of the task of learning Hungarian is to memorize the harmonic behavior (front, back, or vacillating) of the stems that fall into the unpredictable categories (cf. (7) above).

[16] We used OTSoft 2.1 (Hayes, Tesar and Zuraw 2004) to perform the calculation. The grammar was tested ten times, running it 100,000 times for all forms during each trial. The greatest standard deviation across trials for the proportion of back suffixes derived was never greater than .0014 for any input form.

There are various forms of representation that could be used by speakers to memorize whether a stem takes front or back suffixes. These include diacritics, right-aligned floating backness autosegments (Goldsmith 1979), or (following Zuraw 2000) simply the full lexical listing of the inflected forms. These possibilities are shown for the stems in (7) in chart (34).

(34)

<table>
<tr><th></th><th>a. Takes front suffixes</th><th>b. Vacillator</th><th>c. Takes back suffixes</th></tr>
<tr><td>Example</td><td>mutagén</td><td>arzén</td><td>pallér</td></tr>
<tr><td>I. Diacritic</td><td>/mutagén/
[–back harmony]</td><td>/ arzén /
[0back harmony]</td><td>/ pallér/
[+back harmony]</td></tr>
<tr><td>II. Floating autosegment</td><td>/m u t a g é n/
| | |
+b +b –b –b</td><td>/a r z é n/
| |
+b –b –b,

/a r z é n/
| |
+b –b +b</td><td>/p a l l é r/
| |
+b –b +b</td></tr>
<tr><td>III. Full lexical listing</td><td>/mutagén-nek/</td><td>/arzén-nek/,
/arzén-nak/</td><td>/pallér-nak/</td></tr>
</table>

With each of these options, the Faithfulness constraints of the grammar must be stated to enforce the particular form(s) listed in the lexicon. This could be by requiring a proper match between diacritic specification and suffix allomorph (34.I), by requiring surface realization of the floating autosegment (34.II), or simply by requiring the maintenance of underlying suffix vowel backness (34.III).

Here, we adopt the full-listing proposal (34.III), and give some tentative evidence in its favor below. The vowels in listed suffix allomorphs are protected by the general Faithfulness constraint IDENT-IO(back), stated under (24).

Here is an example of how this works. The stem *acél* 'steel' falls into the Bé class, which in the grammar developed so far permits variation in suffix choice. In fact, it is a lexical property of *acél* that it takes only back suffixes. Thus, there is a listed entry /acél-nak/ that emerges as the winner as shown in tableau (35):

(35)

.624

/acél-nak/	IDENT-IO(back)$_{root}$	LOCAL B	IDENT-IO(back)	LOCAL é	DISTAL B
☞ acél-nak		*		*	*
*acél-nek		*	*!		**
*acál-nak	*!		*		

The crucial ranking is IDENT-IO(back) >> LOCAL é, which cancels the possibility that LOCAL é could force the outcome **acél-nek*. Were it not for IDENT-IO(back), this candidate would win 62.4% of the time. The candidate with stem-internal harmony, **acál-nak*, is ruled out by undominated IDENT-IO(back)$_{root}$.[17]

The outcome for *acél* should be compared to the phonologically similar wug stem *hádél*. Wug stems lack lexical entries for their suffixed forms, because the subjects heard only the unaffixed stems. Thus, they cannot specify whether they take -*nak* or -*nek*. Because of this, no candidate violates IDENT-IO(back), and this constraint therefore would not affect the outcome in such cases.

(36)

.624

/hádél-nAk/	IDENT-IO (back)$_{root}$	LOCAL B	IDENT-IO(back)	LOCAL é	DISTAL B
☞ (.624) hádél-nek		*			**
hádél-nak		*		*!	*
hádál-nak	*!		*		

It can be seen that this grammar respects a memorized suffix choice when there is one, but performs stochastically (i.e., like a Hungarian speaker) when given a wug stem. To cover the full range of cases, the rankings needed are as in (37):

(37) IDENT-IO(back) >> {DISTAL B, LOCAL i, LOCAL é, LOCAL e, LOCAL NN}

That is to say, the bloc of stochastically-ranked AGREE constraints in (31) is generally subordinated to IDENT-IO(back).[18]

[17] Zuraw's theory further assumes a constraint USE LISTED, which requires a listed entry to be employed, thus blocking the possibility of a winning candidate **acélnek*, created afresh by the morphology. In the present analysis, USE LISTED may be assumed to be undominated.

[18] Hungarian speakers often feel that "wrong" choices among BN and BNN words (say, *acél-nek*) are not crashingly bad, and that they might accept them if heard from other speakers. This suggests that the main ranking in (37) might actually not be completely strict; and therefore lets through the "wrong" choice as a weak (improbable) alternative. We lack the data that would be needed to establish this stochastic ranking precisely, and will for the remainder of the article assume strict ranking for purposes of exposition.

5.9 The representation of vacillators

For vacillating stems, we follow Ringen and Vago (1995) in assuming that there are two rival underlying representations, as for example (34b) /arzén-nak, arzén-nek/. This is in principle no different from cases like English *envelope*, where /ˈɛnvəˌloʊp/ and /ˈɑnvəˌloʊp/ often compete even within single idiolects. It is likely that the rival underlying forms are represented in a way that assigns them a quantitative "strength", which is reflected by their frequencies in actual usage. The observed Google frequencies of such doublets are seldom actually 50-50 (as the simplest double-listing approach would predict), but vary over the full range of values between 0-100 and 100-0.

This provides an argument for Zuraw's pure lexical-listing theory for idiosyncratic forms ((34.III)). Unlike the diacritic and floating-segment theories, Zuraw's account implies the possibility there could be stems that favor a *different mix* of front and back allomorphs for (say) the dative than for some other suffix. Some supporting cases are given in Kontra and Ringen (1986, 10); and a particularly dramatic case is the common NN stem *férfi* 'man', which is a vacillator in the dative (*-nak/-nek*) but allows only back *-ak* in the plural.[19] In Zuraw's theory, such differences would follow from particular suffixed lexical entries for individual inflected forms. The overall tendency for a stem to take the same backness of suffixes throughout its paradigm would best be attributed to output-to-output correspondence constraints (Benua 1997) governing suffix backness, though we will not attempt to flesh out this proposal here.

5.10 Keeping lexical entries in check

The constraint IDENT-IO(back) permits individual stems to force particular suffix choices, even in the face of the phonological agreement constraints. Yet such suffix preferences should not be allowed unchecked, because the resulting grammar would overgenerate. In Hungarian, there are absolutely no stems of the following types:

(38) *Impossible forms*

a. F stems that take back suffixes (e.g., **gyülüt-nak*)
b. B stems that take front suffixes (e.g., **bortog-nek*)
c. F(N)* stems that take back suffixes (e.g., **gyülit-nak*)

Native speakers vigorously reject such forms, and (other than the occasional random error) do not volunteer them on a wug test.

The normal approach for excluding impossible forms in Optimality Theory follows the principle of the Rich Base (Prince and Smolensky 1993/2004): we show that *if* such items occurred as lexical entries, then the grammar would derive from them an unfaithful well-formed output. For example, if there were a lexical entry like /gyülüt-nak/, the grammar would output *gyülüt-nek* instead.

[19] We owe this observation to Péter Siptár, p.c.

In the grammar under discussion, this can be accomplished by ranking the strongest vowel harmony constraints above IDENT-IO(back), with a probability of 1:

(39)

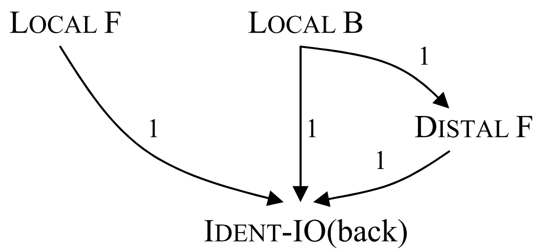

Under this ranking, for the hypothetical lexical entry /gyülüt-nak/ the winning candidate would be well-formed *gyülüt-nek*, in which the underlying /a/ of the suffix surfaces as front:

(40)

/gyülüt-nak/	LOCAL F	IDENT-IO(back)
☞ gyülüt-nek		*
gyülüt-nak	*	

Similar impossible forms are ruled out analogously: because LOCAL B strictly dominates IDENT-IO(back), there could be no forms like **bortog-nek* even if the lexicon "asked for" them (*bortog-nak* wins); and because DISTAL F strictly dominates IDENT-IO(back), there could be no **gyülit-nak* (*gyülit-nek* wins). However, for all Markedness constraints ranked below IDENT-IO(back) (see (37)), an invariant listed form violating that constraint *can* assert itself in the output.

5.11 Summary and assessment

The complete set of constraints and rankings is summarized in the following Hasse diagram:

(41)

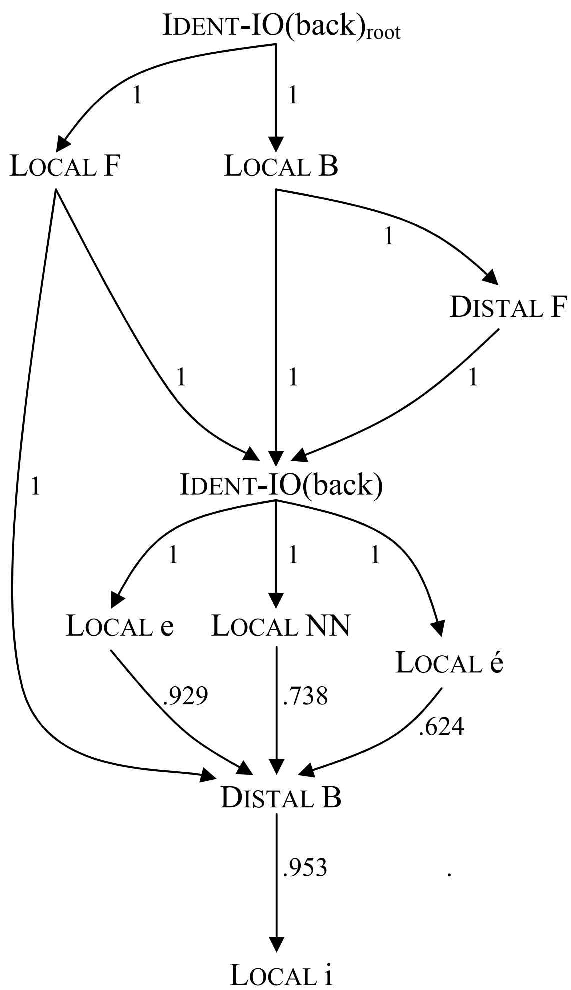

We claim that the analysis meets the goals set out at the beginning of this section. For novel BN and BNN stems, it generates outputs in proportions that fairly closely match those produced by native speakers, as a result of stochastic rankings among the lowest-ranked constraints (§5.7). Further, it permits lexical entries of stems of these types to specify particular suffix choices, as a result of IDENT-IO(back) (§5.8). Lexical entries are not permitted to specify impossible forms, due to the ranking of LOCAL F, LOCAL B, and DISTAL F above IDENT-IO(back) (§5.10). Lastly, because IDENT-IO(back)$_{root}$ is at the top of the grammar, harmony cannot alter stems (§5.6).

The model achieves a fairly close match to the wug test data, diverging slightly in the BNN forms:

(42) *Matchup of model to Wug test data*

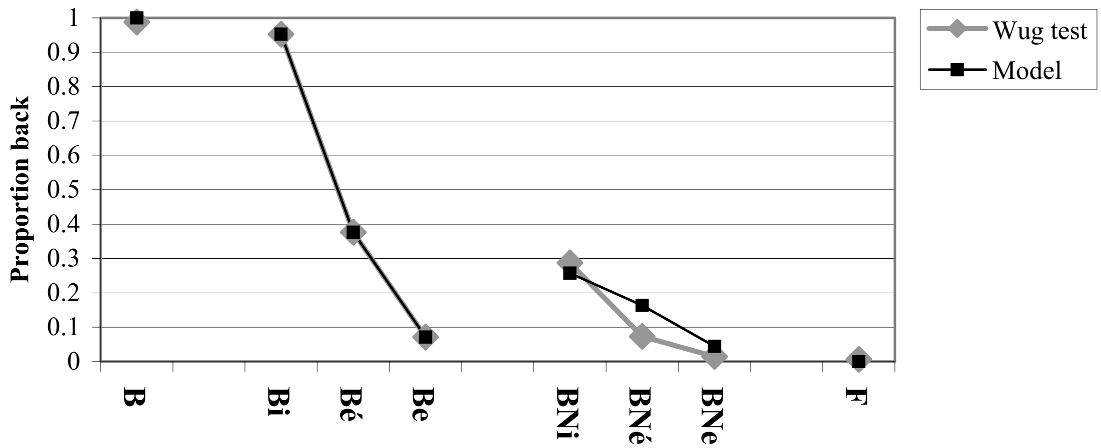

We computed the correlation of the model's predictions against the wug test results across the 30 stems tested (see (14)), and found a value of r = .990, which indicates a fairly close match. This is not surprising given the fairly rich constraint set used. More important, we claim that the constraints themselves are not arbitrary, but follow general principles of phonological theory—that is, they are all either Markedness constraints of the AGREE family or Faithfulness constraints for backness.

5.12 Further issues

A number of issues surrounding the analysis remain for future research.

First, the use of constraints like DISTAL F and DISTAL B glosses over a problem in the analysis of nonlocal harmony in the AGREE framework: as stated, they fail to distinguish forms like BFN from FBN, which incur the same violations. Such forms are rare in Hungarian, but it seems fairly clear that BFN stems take front harmony and FBN stems behave like the corresponding BN stems. In other words, the principle "closest trigger wins" is not predicted by our analysis in the general case.

In the older autosegmental approach, the "closest trigger wins" principle was the automatic consequence of the ban on crossed association lines (Clements 1976). However, in light of recent research (Hansson 2001; Frisch, Pierrehumbert, and Broe 2004; Rose and Walker 2004) it appears that the "closest trigger wins" principle is not exceptionless; and it is entirely incompatible with the view taken here, in which variation in Hungarian is attributed to conflicting harmony triggers. What is needed, we think, is a more articulated theory of nonlocal AGREE constraints in which the principle becomes negotiable, subject to constraint ranking.

One possibility for such a theory is to suppose that a constraint like DISTAL B is actually an intrinsically ranked infinite schema, along these lines:

(43) AGREE B ___ >>
AGREE BX ___ >>
AGREE BXX ___ >>
AGREE BXXX ___ >>
…

What needs to be worked out is a constrained system for ranking the members of this schema which would appropriately implement the general idea of "closest trigger wins." For instance, we expect the default ranking of the DISTAL B schema and the DISTAL F schema to be as in (44):

(44)

AGREE B ___	AGREE F ___
AGREE BX ___	AGREE FX ___
AGREE BXX ___	AGREE FXX ___
AGREE BXXX ___	AGREE FXXX ___
…	…

Since developing such a system explicitly would lead us rather far from present concerns we will not attempt it here.

We also have not dealt with the N and NN stems. In our overall approach, the occasional volunteering of back suffixes for N stems in our wug test (as in *híny-nak*; see §4.2) cannot be the result of IDENT-IO(back), since wug stems are assumed to lack lexical entries. Rather, they must reflect constraints, perhaps language-particular, that require dissimilation (cf. Ringen 1980, 139). Most likely, the principal dissimilation constraint singles out monosyllables in /í/, since most of the *híd* stems of Hungarian take this form. The dissimilation constraints must be ranked stochastically, lower than the constraints of the LOCAL N family, so that wug forms attaching *-nak* to neutral stems are output only occasionally.

6. Learning the Grammar

Thus far, we have followed the classical procedure of generative linguistics, inventing a formal hypothesis intended to describe the native speaker's tacit knowledge. But the goal of explanatory adequacy (Chomsky 1965) implies that linguistic theory should not just provide an accurate and principled account of language-particular grammars, but offer a mechanism for how a child exposed to learning data could arrive at the correct grammar. Recent work on learnability in Optimality Theory (Boersma 1997, Boersma and Hayes 2001, Hayes 2004, Prince and Tesar 2004) makes possible a sketch of how the Hungarian vowel harmony system as outlined here could be learned. The discussion will be confined to the learning of the rankings shown in (41); we assume that the constraint inventory itself is either innate (Tesar and Smolensky 2000) or else is accessible to the child through some form of inductive learning (Hayes 1999).

6.1 Factoring the learning task

The learning task at hand can be divided into three parts.

(a) The child must discover **what is possible**; that is, she learns to distinguish harmonically possible words from harmonically impossible ones. For instance, she must ultimately come to know that forms such as **gyülüt-nak*, **bortog-nek*, **gyülit-nak* (all from (38)), **glükóz-nek* ((26)), and **sofőr-nak* ((27)) are all impossible in Hungarian; and must also learn that forms like *hádél-nak* and *hádél-nek* are both, in principle, possible Hungarian forms.

(b) Another aspect of the learning task involves the **lexicon**: the child must learn which particular stems take which kinds of suffixes, internalizing lexical entries along the lines of §5.8.[20]

(c) Lastly, it is evident from our experimental results that Hungarian-learning children ultimately develop a **statistical model** of the lexicon, hence the ability to project novel forms stochastically in proportions matching their lexical frequencies. In our model, this is done by assigning a stochastic ranking to the constraints at the bottom of the grammar in (41).

We will ignore here the (non-trivial) task of learning thousands of lexical entries and focus instead on tasks (a) and (c). In principle, they could be accomplished by a single algorithm, but here we find it necessary to use two. The scheme invoked here is to use one algorithm to learn what is possible, by establishing a set of ranked constraint strata, plus a second algorithm to fine tune these strata with statistical information. The two algorithms could in principle run simultaneously; what is crucial is that the statistical fine-tuning must respect the overall stratal structure.

6.2 Learning what is legal

The task of learning what is possible in Hungarian vowel harmony confronts a classical conundrum: no negative evidence is available. The child is never informed that words like **gyülüt-nak*, **bortog-nek*, etc., are ill-formed, but rather comes to know it through some combination of Universal Grammar and data processing capable of detecting the systematic gaps in the learning data.

For Optimality Theory, this kind of problem has been addressed with two proposed constraint ranking algorithms: Low Faithfulness Constraint Demotion (Hayes 2004) and Biased Constraint Demotion (Tesar and Prince 2004). In our learning simulations, we tried both of these algorithms.

We fed the algorithms data that specified the kinds of forms that exist in Hungarian, but without any frequency information. The data were schematic forms like "B-*nak*," intended to

[20] Until the child learns harmony, these lexical entries might include even completely predictable cases like *ablak-nak* 'window-dat.' Once the child knows that B stems always take back suffixes, such entries could be allowed to atrophy (or not; cf. Baayen et al. 2002 and literature cited there). For unpredictable cases like *acél-nak* (cf. (35)), the memorized entries necessarily must be retained into adulthood.

express whole classes of real Hungarian stems that share the same constraint violations. The learning data were as follows:

(45)

a. B-*nak*	c. Bi-*nak* Bi-*nek*	f. BNi-*nak* BNi-*nek*	i. FB-*nak*	l. BF FB FNB BNF
b. F-*nek*	d. Bé-*nak* Bé-*nek*	g. BNé-*nak* BNé-*nek*	j. BF-*nek*	
	e. Be-*nak* Be-*nek*	h. BNe-*nek*	k. Fi-*nek*	

The presence of both Bi-*nak* and Bi-*nek* in the learning data (see cell (45c)) meant that a ranking permitting both types had to be found; this turns out to be Ident-IO(back) >> {Distal B, Local i}; which can be seen in (41) above. The fact that B-*nek* is *not* in the learning data (cell (45a)) means that a ranking must be found that excludes it; this turns out to be Local B >> Ident-IO(back). Cells (45b, i-k) similarly support the rankings in (41) that result in these gaps. The fact that FB is in the learning data (cell (45l)) means that a ranking must be discovered that permits it; this turns out to be Ident-IO(back)$_{\text{root}}$ >> Local F; and similarly for the other forms in the same cell.

The form BNe-*nak* is not in the learning data (see (45h)); as noted above, such forms do not occur in Hungarian. However, if we are to account for our wug test data, where forms like BNe-*nak* were actually volunteered by native speakers, the learned grammar must be able to generate it, even though it is not part of the learning data.

Both of our algorithms, Low Faithfulness Constraint Demotion and Biased Constraint Demotion, require a set of losing candidates for learning (the rankings are learned by comparing these losers with winning candidates). We obtained our losing candidates using the method given in Tesar and Smolensky (2000) and Tesar and Prince (2004): they are simply the wrong guesses made by preliminary versions of the grammar.

We will not review the specific courses followed by the algorithms, but simply give the results they obtained. The two algorithms learned identical, correct grammars, in the form of the strictly ranked constraint strata given below:

(46) Ident-IO(back)$_{\text{root}}$ >>
{ Local B, Local F } >>
Distal F >>
Ident-IO(back) >>
{ Distal B, Local i, Local é, Local e, Local NN }

From (46), all the nonstochastic rankings (probability = 1) of Hasse diagram (41) can be deduced. As already shown, these rankings guarantee that none of the impossible forms mentioned above can be generated, and all of the possible forms can.[21]

[21] For the reason just given, we consider BNe-*nak* to be a possible form.

6.3 Learning statistical distributions

The other part of learning consisted of fine-tuning this overall ranking so as to match the frequencies of the lexicon. For this purpose, we used the Gradual Learning Algorithm (Boersma 1997, Boersma and Hayes 2001), operating under the constraint that it had to respect all of the pairwise rankings defined by the strata in (46). Under this regimen, most stem types cannot influence ranking, so we (harmlessly) restricted the learning set to the stem types that matter, namely Bi, Bé, Be, BNi, BNé, and BNe. As a means of approximating the experience of real Hungarian learners, we chose as frequencies the type frequencies found in the Google survey reported in §3.1, dividing the frequency share of each vacillator in proportion to the token counts. Thus the learning data assumed were as follows:

(47)

Stem type	*Harmony pattern*	*Frequency*	*Stem type*	*Harmony pattern*	*Frequency*
Bi	Bi-nak	469.8	BNi	BNi-nak	6.7
	Bi-nek	5.2		BNi-nek	23.3
Bé	Bé-nak	101.4	BNé	BNé-nak	5.1
	Bé-nek	18.6		BNé-nek	6.9
Be	Be-nak	12.1	BNe	BNe-nak	0.0
	Be-nek	103.9		BNe-nek	21.0

We further assumed that this part of learning is uninfluenced by Faithfulness, in particular Faithfulness to listed suffixed forms like (34c) /pallér-nak/. This assumption was needed to keep Faithfulness from determining the outcome, which would have kept the GLA from learning the aggregate statistical pattern of the lexicon. We therefore excluded the Faithfulness constraints from this phase of the ranking; they are all ranked correctly in any event in the non-stochastic phase just described.[22]

We ran the Gradual Learning Algorithm for ten trials on these learning data in the way just described. All trials yielded similar outcomes; we report the least accurate one here.[23]

For the stochastic lower region of the grammar, which is what is at issue here, the algorithm learned the stochastic rankings given below. We give pairwise ranking probabilities for the four crucial cases, where DISTAL B, which favors back suffixes, conflicts with some other constraint favoring front suffixes:

---

[22] A more nuanced approach would suppose that the irrelevance of IDENT-IO(back) arises not from simply turning Faithfulness off, but rather from the fact that it takes time for lexical entries like /pallér-nak/ to get established—it would take multiple hearings for a learner to become confident that *pallér* is not a vacillator. During this period, when the learner hears *pallér-nak*, she can only interpret it as a representative Bé stem. In this capacity, it would play a role in incrementally reranking the Markedness constraints to favor back suffixes for such stems.

[23] The learning parameters were: 250,000 trials at each of the plasticity values 1, .1, .01, and .001; noise set at 2.0 for all trials; results tested for 100,000 trials. Strict rankings were enforced by maintaining a minimum distance of 20 along the ranking scale. The ranking values output by the GLA were LOCAL e = 105.176, LOCAL NN = 103.313, DISTAL B = 101.430, LOCAL é = 98.315, LOCAL i = 95.079.

(48) *Ranking* *Probability*

a.	LOCAL e >> DISTAL B	.903
b.	LOCAL NN >> DISTAL B	.739
c.	DISTAL B >> LOCAL é	.871
d.	DISTAL B >> LOCAL i	.988

6.4 Evaluating the simulation

The learned grammar can be evaluated in two ways. First, we can ask the purely mechanical question of whether the algorithm was able to rank the constraints in a way that mimicked the frequencies (that is, the Google frequencies) in the learning data. From a more scientific viewpoint, we can ask if it mimicked a real Hungarian speaker: ideally, the learning system should be trained with real data, then behave like a native speaker when it is wug tested—including any divergences from the pattern in the learning data.

The results for the criterion of corpus-mimicry were reasonably good (correlation for the eight values given: r = .988) and are given in the following chart.

(49) *Matchup of machine-ranked model to Google data*

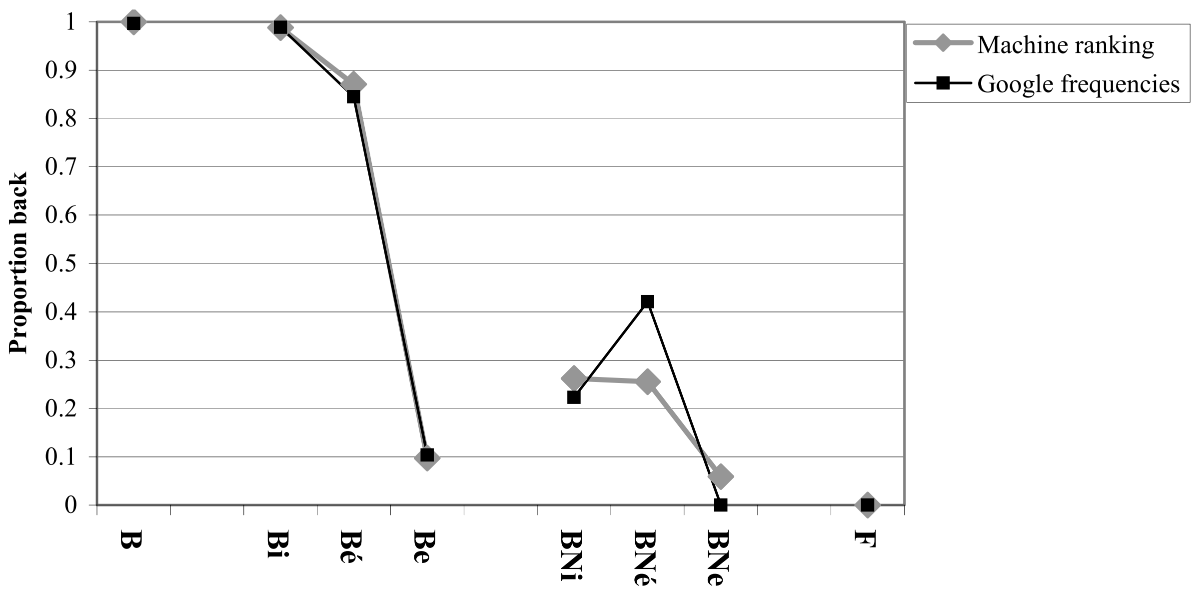

It is clear that the B and F forms were unproblematic, and that the algorithm was driven primarily by the frequencies of the BN forms, which were closely mimicked. The frequencies of BNN forms were not so accurately reflected; however, we would claim that this is all to the good: like native speakers, the algorithm *smoothed* these forms (§4.3; §5.7). The peak in BNé forms is diminished in the learned grammar, and the fraction of BNe-*nak* forms is raised from zero to a modest level. These smoothings are the result of the algorithm carrying over patterns from the statistically preponderant BN forms.

Next, we compare the results of the learning model to the wug test data. The matchup achieved by our model is in (50) below.[24] For comparison we include as well the predictions made by the handcrafted grammar of (41) (i.e., repeating the rows from chart (42)).

(50) *Matchup of models (hand-ranked, machine-ranked) to Wug test data*

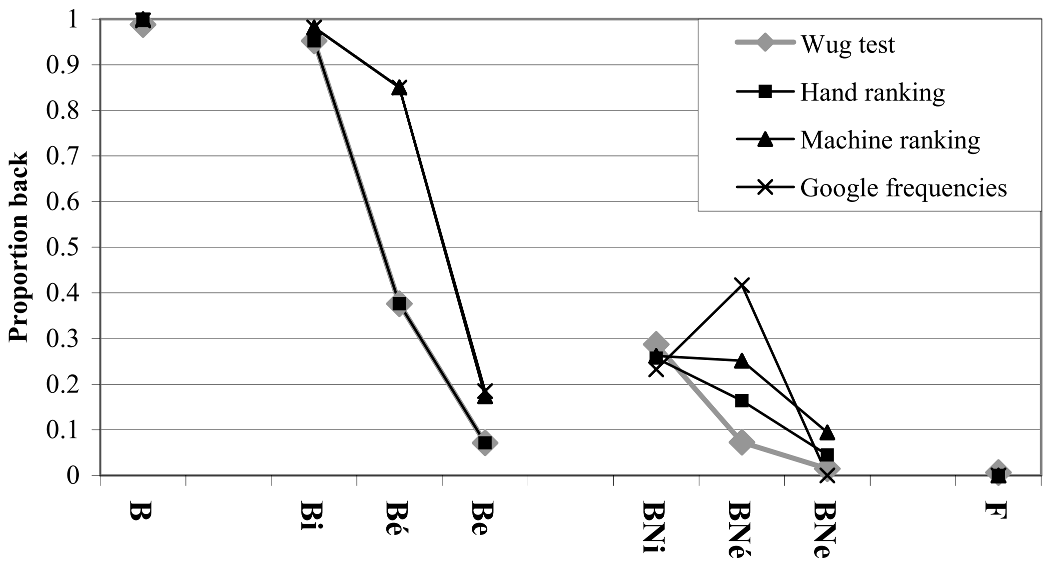

The following observations seem pertinent:

(a) The machine-learned grammar did somewhat worse than the handcrafted grammar at matching the wug test data (correlations: r = .908 vs. r = .990). This is to be expected, since the handcrafted grammar was deliberately made to match the wug test data, whereas the machine-learned grammar was trained on the Google data.[25]

---

[24] These are the numerical values in chart (50):

Stem Type	Wug test	Hand ranking	Machine ranking	Google frequencies
B	0.988	1.000	1.000	0.997
Bi	0.953	0.953	0.988	0.989
Bé	0.376	0.376	0.871	0.845
Be	0.071	0.071	0.098	0.104
BNi	0.287	0.257	0.262	0.223
BNé	0.073	0.164	0.255	0.421
BNe	0.015	0.045	0.059	0.000
F	0.006	0.000	0.000	0.000

[25] Note that running a learning simulation in which the learning data are wug test intuitions would lack scientific legitimacy; it would presuppose that language learners could directly access the intuitions of those around them rather than just their speech output.

(b) As noted, the machine-learned grammar smoothed the frequencies of BNN forms. As a result, the correlation of the wug test data (30 forms) with the predictions of the machine-learned grammar ($r = .908$) was actually slightly higher than the correlation of the wug test data with Google data from which the machine grammar was learned (among the forms covered by the grammar, this is $r = .896$). In other words, the machine-learned grammar's failures in mimicking its learning data actually made it a slightly better mimic of the wug test data, a desirable outcome given our goal of modeling real language learners.

(c) However, the machine-learned grammar evidently did not smooth as much as it should have. The handcrafted grammar, which was set up to match the wug test data rather than the Google data, achieves a better degree of smoothing for BNé (lower value) than the machine-learned grammar.

(d) A further source of discrepancy between the machine-learned model and the wug test data is the apparent bias for front harmony among the wug testees relative to the lexicon, already discussed in §4.4.

Overall, we are encouraged by the degree of match between the learning model and the wug test data, and by the ability of the model to smooth in a qualitatively appropriate way.

7. Conclusions

Our main empirical result, from the wug test, is that Hungarian speakers know not just the legal patterns of harmony, but also the frequency of these patterns, and they actively use this knowledge in guessing the harmonic behavior of novel stems. This is linguistic knowledge that most previous models of phonology and morphology (see §1) have not captured.

It is not implausible to suppose that this knowledge is in fact useful to speakers and thus worth acquiring. It permits them to guess more accurately when they must produce an inflected form for a stem they have never heard with a suffix—probably a common experience for children. It may also serve them in speech perception, by providing rational top-down biases for recognizing suffixes uttered by other speakers.

On the theoretical side, we have suggested that the way Hungarian speakers internalize the frequencies is not through some kind of raw data table, but instead in their grammars. We have found that by ranking some of the constraints stochastically in a "subterranean" grammar of the kind proposed by Zuraw (2000), we can model the native speaker's intuition fairly accurately. The constraints that are needed are ordinary constraints of Optimality Theory; all that is different is the possibility of stochastic ranking. The fact that stochastic OT also allows rankings that are essentially non-stochastic (probability vanishingly close to 1) means that our model can also rule out impossible forms.

Lastly, we have made an initial attack on the problem of learning such grammars. Our conjecture is that the full set of rankings can be learned through a combination of algorithms, one of which learns the basic range of possibilities, while the other fine-tunes the grammar to match lexical frequencies.

A theme of our work has been to show that tools now exist to permit study of phonological systems in greater detail than would otherwise be possible. Our web-corpus study, experimentation, and stochastic theoretical modeling show that the native speaker of Hungarian possesses a richer knowledge of the harmony system than can be adequately described under the older rules-and-exceptions approach.

This said, we believe that a great deal of further progress is needed. In particular, while we have demonstrated the considerable detail with which the native speaker of Hungarian learns the harmony pattern of the lexicon, our study was not designed to find out any upper limit on what is learned: are there statistically reliable patterns in the lexicon that cannot be detected by the human phonological capacity?[26] A positive answer to this question would be very informative concerning the nature of that capacity. We think the methods laid out here might serve to address this question, notably by expanding greatly (perhaps by use of the Web) the scope of wug testing. We also anticipate that further progress will follow from improvements in the underlying phonological theory and in the theory of phonological learning.

References

Albright, Adam (2002). Islands of reliability for regular morphology: Evidence from Italian. *Language* **78**. 684-709.

Albright, Adam and Bruce Hayes (2003). Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition* **90**. 119-161

Anderson, Lloyd (1980). Using asymmetrical and gradient data in the study of vowel harmony. In Robert M. Vago (ed.), *Issues in Vowel Harmony*. Amsterdam: J. Benjamins, 271-340.

Archangeli, Diana and Douglas Pulleyblank (1987). Maximal and minimal rules: effects of tier scansion. In Joyce McDonough and Bernadette Plunkett, eds., *Proceedings of the North Eastern Linguistic Society* 17, Graduate Linguistics Student Association, University of Massachusetts, Amherst, 16-35.

Baayen, Harald, Robert Schreuder, Nivja De Jong, and Andrea Krott (2002). Dutch inflection: The rules that prove the exception. In Sieb Nooteboom, Frank Wijnen and Fred Weerman (eds.), *Storage and computation in the language faculty*. Dordrecht: Kluwer, 61-92.

Benua, Laura (1997). *Tranderivational identity: phonological relations between words*. Ph.D. dissertation, University of Massachusetts, Amherst.

Benus, Stefan (2005). *Dynamics and transparency in vowel harmony*. Ph.D. dissertation, Department of Linguistics, New York University.

Benus, Stefan, and Adamantios Gafos (in press). Qualitative and quantitative aspects of vowel harmony: a dynamics model. To appear in *Proceedings of the 27th Annual Conference of the Cognitive Science Society*.

Berko, Jean (1958). The child's learning of English morphology. *Word* **14**. 150-177.

[26] In the Google data, we find small effects from the height of the penultimate vowel of BNN, and also (as in Finnish, Ringen and Heinämäki 1999) of the main stressed (initial) vowel of the word. Given the few forms that we wug-tested, we lack the data to determine whether these effects are actually internalized by native speakers of Hungarian.

Boersma, Paul (1997). How we learn variation, optionality, and probability. *Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam* **21**. 43–58.

Boersma, Paul and Bruce Hayes (2001). Empirical tests of the Gradual Learning Algorithm. *Linguistic Inquiry* **32**. 45-86.

Bybee, Joan (1995). Regular morphology and the lexicon. *Language and Cognitive Processes* **10**. 425–455.

Bybee, Joan (2001). *Phonology and language use*. Cambridge: Cambridge University Press.

Bybee, Joan and Carol Moder (1983). Morphological classes as natural categories. *Language* **59**. 251-270.

Casali, Roderic (1997). Vowel elision in hiatus contexts: which vowel goes? *Language* **73**. 493-533

Chomsky, Noam (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Clements, George N. (1976). Neutral vowels in Hungarian vowel harmony: an autosegmental interpretation. In Judy Kegl, David Nash, and Annie Zaenen, eds., *Proceedings of the Seventh Annual Meeting of North Eastern Linguistic Society*, 49-64.

Clements, George N. and Elizabeth V. Hume (1995). The internal organization of speech sounds. In John Goldsmith, ed., *The Handbook of Phonological Theory*. Oxford: Blackwell, 245-306.

De Lacy, Paul (2004). Markedness conflation in Optimality Theory. *Phonology* **21**. 145-199.

Dienes, Péter (1997). Hungarian neutral vowels. *The Odd Yearbook*, 151-180.

Eddington, David (1996). Diphthongization in Spanish derivational morphology: an empirical investigation. *Hispanic Linguistics* **8**. 1-35

Eddington, David (2004). *Spanish Phonology and Morphology: Experimental and Quantitative Perspectives*. Amsterdam: John Benjamins.

Ernestus, Miriam and Harald Baayen (2003). Predicting the unpredictable: Interpreting neutralized segments in Dutch. *Language* **79**. 5-38.

Esztergár, Marianne (1971). *A generative phonology of nouns and vowel harmony in Hungarian*. Ph.D. dissertation, University of California, San Diego.

Farkas, Donka and Patrice Beddor (1987). Private and equipollent backness in Hungarian. In Anna Bosch, Barbara Need, and Eric Schiller, eds., *Papers from the Twenty-Third Annual Regional Meeting of the Chicago Linguistic Society, Part II: Parasession on Autosegmental and Metrical Phonology*. Chicago: Chicago Linguistic Society, 91-105.

Frisch, Stefan, Janet Pierrehumbert, and Michael Broe (2004). Similarity avoidance and the OCP. *Natural Language and Linguistic Theory* **22**. 179-228.

Goldsmith, John (1979). *Autosegmental Phonology*. New York: Garland.

Gósy, Mária (1989). Vowel harmony: interrelations of speech production, speech perception, and the phonological rules. *Acta Linguistica Hungarica* **39**. 98-118.

Halle, Morris and K. P. Mohanan (1985). Segmental phonology of modern English. *Linguistic Inquiry* **16**. 57-116.

Hansson, Gunnar (2001). *Theoretical and Typological Issues in Consonant Harmony*. Ph.D. dissertation, UC Berkeley.

Hare, Mary (1990). The role of similarity in Hungarian vowel harmony: A connectionist account. *Connection Science*, **2**.123-150.

Hayes, Bruce (1999). Phonetically-driven phonology: the role of optimality theory and inductive grounding. In Michael Darnell, Edith Moravcsik, Michael Noonan, Frederick Newmeyer, and Kathleen Wheatly, eds., *Functionalism and Formalism in Linguistics, Volume I: General Papers.* Amsterdam: John Benjamins, 243-285.

Hayes, Bruce (2004). Phonological acquisition in Optimality Theory: the early stages. In René Kager, Joe Pater, and Wim Zonneveld (eds.), *Fixing priorities: constraints in phonological acquisition*. Cambridge University Press, 158-203.

Hayes, Bruce and Margaret MacEachern (1998). Quatrain form in English folk verse. *Language* **74**. 473-507.

Hayes, Bruce, Bruce Tesar, and Kie Zuraw (2003). OTSoft 2.1. Software package, http://www.linguistics.ucla.edu/people/hayes/otsoft/.

Kaun, Abigail (1995). *The Typology of Rounding Harmony: An Optimality Theoretic Approach*. Ph.D. dissertation, UCLA.

Kaun, Abigail (2004). The typology of rounding harmony. In Hayes, Bruce, Robert Kirchner, and Donca Steriade, eds., *Phonetically-Based Phonology*. Cambridge: Cambridge University Press, 87-116.

Kiparsky, Paul and Karl Pajusalu (2003). Towards a typology of disharmony. *The Linguistic Review* **20**. 217-241.

Kontra, Miklós and Catherine Ringen (1986). Hungarian vowel harmony: the evidence from loanwords. *Ural-Altaische Jahrbücher* **58**.1-14

Kornai, András (1991). Hungarian vowel harmony. In István Kenesei (ed.) *Approaches to Hungarian III*, 183-240

Lombardi, Linda (1999). Positional faithfulness and voicing assimilation in optimality theory. *Natural Language and Linguistic Theory* **17**. 267-302.

McCarthy, John and Matthew Wolf (2005). Less than zero: Correspondence and the null output. Rutgers Optimality Archive 722, http://roa.rutgers.edu.

Pierrehumbert, Janet (2001). Stochastic phonology. *GLOT* **5**. 1-13.

Pierrehumbert, Janet (in press). An unnatural process. To appear in Louis Goldstein, ed., *Papers in Laboratory Phonology VIII*. Berlin: Mouton de Gruyter.

Pinker, Steven and Alan S. Prince (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition* **28**. 73–193.

Prasada, Sandeep and Steven Pinker (1993). Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes* **8**. 1–56.

Prince, Alan (1997). Stringency and anti-Paninian hierarchies. Talk given at Cornell University. http://ling.rutgers.edu/people/faculty/prince.html.

Prince, Alan and Paul Smolensky (1993/2004). *Optimality Theory: Constraint Interaction in Generative Grammar*. Oxford: Blackwell. [1993 version: www.roa.rutgers.edu].

Prince, Alan and Bruce Tesar (2004). Learning phonotactic distributions. In R. Kager, J. Pater, and W. Zonneveld (eds.), *Fixing priorities: constraints in phonological acquisition*. Cambridge University Press, 245-291.

Reiss, Charles (2003). Deriving the feature-filling/feature-changing contrast: an application to Hungarian vowel harmony. *Linguistic Inquiry* **34**. 199-224.

Ringen, Catherine O. (1975). *Vowel Harmony: Theoretical Implications*. Ph.D. dissertation, Indiana University. Published 1988: New York: Garland.

Ringen, Catherine O. (1980). A concrete analysis of Hungarian vowel harmony. In Robert M. Vago, ed., *Issues in Vowel Harmony*. Amsterdam: John Benjamins, 134-154.

Ringen, Catherine O. and Orvokki Heinämäki (1999). Variation in Finnish vowel harmony. *Natural Language and Linguistic Theory* **17**. 303–37.

Ringen, Catherine O. and Robert M. Vago (1995). A constraint based analysis of Hungarian vowel harmony. In István Kenesei (ed.) *Approaches to Hungarian, Vol. 5: Levels and structures*. Szeged: JATE, 309-319.

Ringen, Catherine O. and Robert M. Vago (1998). Hungarian vowel harmony in Optimality Theory. *Phonology* **15**. 393-416.

Rose, Sharon and Rachel Walker. (2004). A typology of consonant agreement as correspondence. *Language* **80**. 475-531.

Siptár, Péter and Miklós Törkenczy (2000). *The Phonology of Hungarian*. Oxford: Oxford University Press.

Szépe, György (1958). Vegyes magánhangzójú szavaink illeszkedésének kérdéséhez. In Benkő, Loránd (ed.), *Magyar hangtani dolgozatok* (= *Nyelvtudományi Értekezések* 17). Budapest: Akadémiai Kiadó, 105–129.

Tesar, Bruce and Paul Smolensky (2000). *Learnability in Optimality Theory*. Cambridge, MA: MIT Press.

Vago, Robert M. (1973). Abstract vowel harmony systems in Uralic and Altaic languages. *Language* **49**. 579-605.

Vago, Robert M. (1974). *Hungarian Generative Phonology*. Ph.D. dissertation, Harvard University.

Vago, Robert M. (1976). Theoretical implications of Hungarian vowel harmony. *Linguistic Inquiry* **7**. 243-263.

Vago, Robert M. (1980). *The Sound Pattern of Hungarian*. Washington: Georgetown University Press.

Vergnaud, Jean-Roger and Morris Halle (1979). Metrical structures in phonology. Ms., Dept. of Linguistics and Philosophy, MIT, Cambridge, MA.

Walker, Rachel (2001). Round licensing, harmony, and bisyllabic triggers in Altaic. *Natural Language and Linguistic Theory* **19**. 827-878.

Zuraw, Kie (2000). *Patterned Exceptions in Phonology*, Ph.D. dissertation, UCLA.